OSH / FARO

Physical RC Mutant Plant Benchmark

Slew-Floor Corrected Analysis

FARO Compacta V2 + OAPA V11 vs fixed R6 PID
CH32V307VCT6 physical RC plant | 29 October 2026

Raw IAE: FARO 114,819.00 vs PID 116,345.88 - 1.31% lower

Above slew-only floor: FARO 17,077.56 vs PID 18,604.44 - 8.21% lower

Methodological correction: the original raw metric is preserved. The new normalization removes only the analytically unavoidable command-slew contribution and is reported as a post-hoc secondary analysis, not as a replacement for the primary result.

Executive summary

The previous consolidated report correctly reported the raw nominal IAE, but that metric is dominated by a protocol-imposed actuator slew limit. This revision preserves the original raw result and adds an exact slew-floor decomposition so the remaining separation is visible without overstating what the normalization proves.

Measure	FARO	PID	Interpretation
Raw nominal IAE	114,819.00	116,345.88	1.31% lower with FARO
Shared slew-only floor	97,741.44	97,741.44	Identical analytical floor across 4 profiles
IAE above slew floor	17,077.56	18,604.44	8.21% lower with FARO
Mean tail RMSE	1.581	5.079	68.87% lower with FARO
Total DAC variation	23,826	20,770	14.71% higher with FARO

Exact correction to the approximate claim. The slew-only floor is **24,435.36 count-s per profile**, not 24,600. Across the four profiles it is 84.01% of PID raw IAE, 85.13% of FARO raw IAE, or 84.56% of the pair mean. The above-floor relative reduction is **8.21% aggregate**, with a per-profile range of **6.12% to 11.20%**. Therefore this report does not state a fixed 86% floor or a 7-13% range.

Important: the above-floor remainder is not labelled “100% controllable error.” The real RC plant has dynamics of its own, so part of that remainder may also be physically unavoidable. The normalization isolates only the guaranteed contribution of the 12-count command slew constraint.

Experiment status: the FAST and SLOW blocks are a descriptive consolidation separated by a contact intervention. Timing comparability passed within the completed pairs, but the stored result explicitly does not authorize treating the combined set as one uninterrupted physical experiment.

1. Physical benchmark protocol

The controller interface operated a physical RC plant through the CH32V307VCT6 DAC/ADC path. The host controllers were FARO Compacta V2 + OAPA V11 with adaptation enabled and a fixed R6 PID. Each scored profile contained 540 decisions at a nominal 120 ms period.

Item	Frozen value
Scored profiles	FAST_UP, FAST_DOWN, SLOW_UP, SLOW_DOWN
Decisions per profile	540
Nominal period	0.120 s
Action limit	12 DAC counts per decision
Primary metric	0.12 x sum(|reference - raw post median ADC|)
Preparation	384 decisions per controller on MEDIUM; no post-mutation training
Timing acceptance	<=5% cycles >150 ms and paired duration ratio <=1.10
Plant mutation	Manual between powered-off blocks; not continuous hot-switch

The scored reference sequences are symmetric in jump magnitude. UP follows 2048 -> 2816 -> 1280 -> 2688; DOWN follows 2048 -> 1280 -> 2816 -> 1408. Each target is held for 180 decisions. The three jump magnitudes are therefore 768, 1536 and 1408 counts.

Because the action limit is applied on every decision, a controller cannot instantaneously move the command to a new setpoint. That shared restriction contributes a large amount of tracking error before controller-specific settling behavior can separate the arms.

2. Exact slew-rate floor

Define D as the absolute reference jump, L = 12 counts/decision, and dt = 0.12 s. Under a best-case instantaneous unity-gain plant and a causal controller with no preview of the next reference, the minimum post-action error after decision k is max(D - kL, 0).

```
B(D) = dt x sum[k=1..ceil(D/L)] max(D - kL, 0)
```

Jump D	Moves to target	Sum of unavoidable count-error	IAE floor
768	64	24,192	2,903.04
1,536	128	97,536	11,704.32
1,408	118	81,900	9,828.00
Per profile		203,628	24,435.36

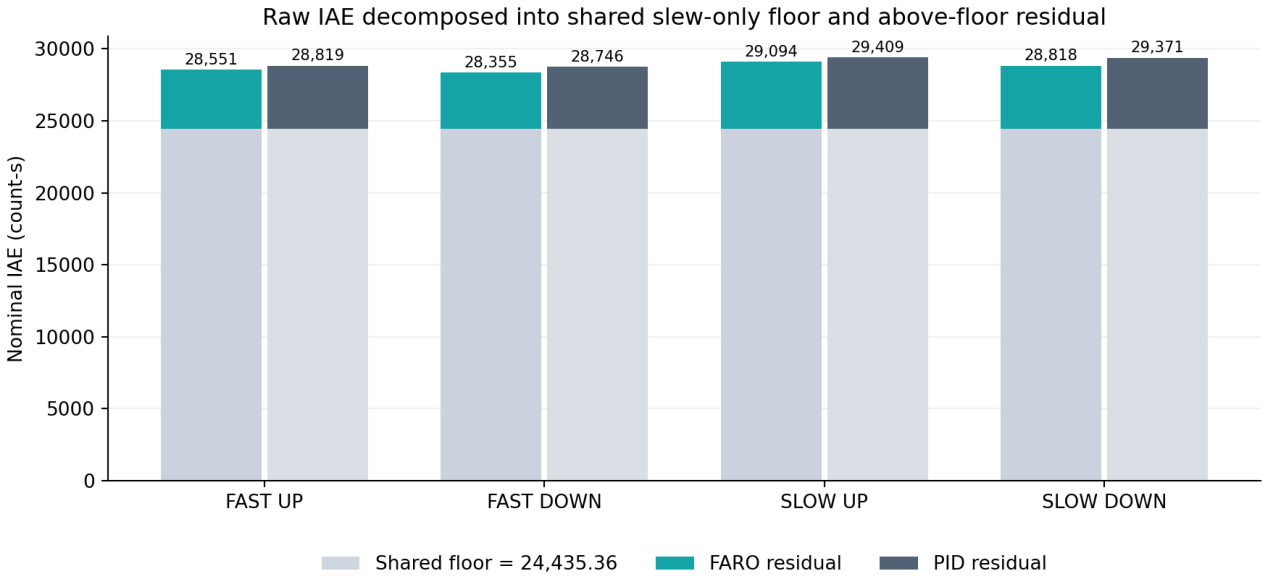
Thus the exact shared slew-only floor is **24,435.36 count-s per profile**. Across four profiles it is **97,741.44** count-s.

This is deliberately a conservative decomposition: it assumes the physical output can follow the DAC command immediately. The actual RC plant cannot do better than that idealization under normal passive first-order behavior, so real plant dynamics can add error above this floor. We therefore call the remainder “IAE above the slew-only floor,” not pure controller-controllable error.

This decomposition was derived after the physical results were known. It is a secondary post-hoc analysis. The preregistered/raw nominal IAE remains reported unchanged.

3. Why the raw IAE looks compressed

Most of the measured nominal IAE is occupied by the same command-ramp burden. The chart separates that common floor from the remainder where plant dynamics and controller behavior can produce separation.

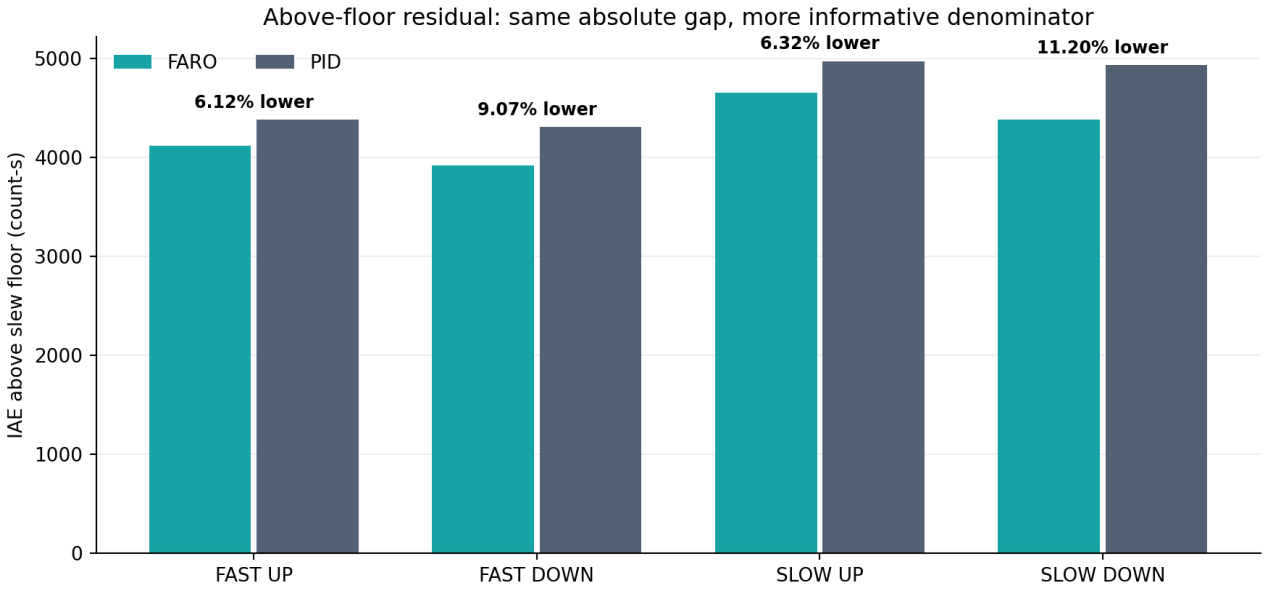


Profile	Floor / PID raw	Floor / FARO raw	Raw FARO reduction
FAST_UP	84.79%	85.58%	0.93%
FAST_DOWN	85.00%	86.18%	1.36%
SLOW_UP	83.09%	83.99%	1.07%
SLOW_DOWN	83.19%	84.79%	1.88%

Using PID as the denominator, the shared floor accounts for **83.09% to 85.00%** of each profile and **84.01% aggregate**. This is the precise version of the “roughly 86%” observation.

4. Above-floor residual result

Subtracting the same analytically derived floor from both arms leaves the absolute FARO-PID gap unchanged, but changes the denominator used to express relative separation.



Profile	FARO residual	PID residual	FARO lower vs PID residual
FAST_UP	4,115.76	4,384.08	6.12%
FAST_DOWN	3,919.92	4,311.00	9.07%
SLOW_UP	4,659.00	4,973.52	6.32%
SLOW_DOWN	4,382.88	4,935.84	11.20%
Aggregate	17,077.56	18,604.44	8.21%

Interpretation: FARO is lower in all four above-floor residuals. The aggregate reduction is 8.21%. This does not mean FARO created an 8.21% reduction in total physical error; the absolute reduction is exactly the same 1,526.88 count-s as in the raw metric. The normalization shows that this gap is a larger fraction of the error remaining after removal of the shared slew-only burden.

5. Raw results remain primary

The original nominal IAE values are retained verbatim. No physical run was rescored and no threshold was changed.

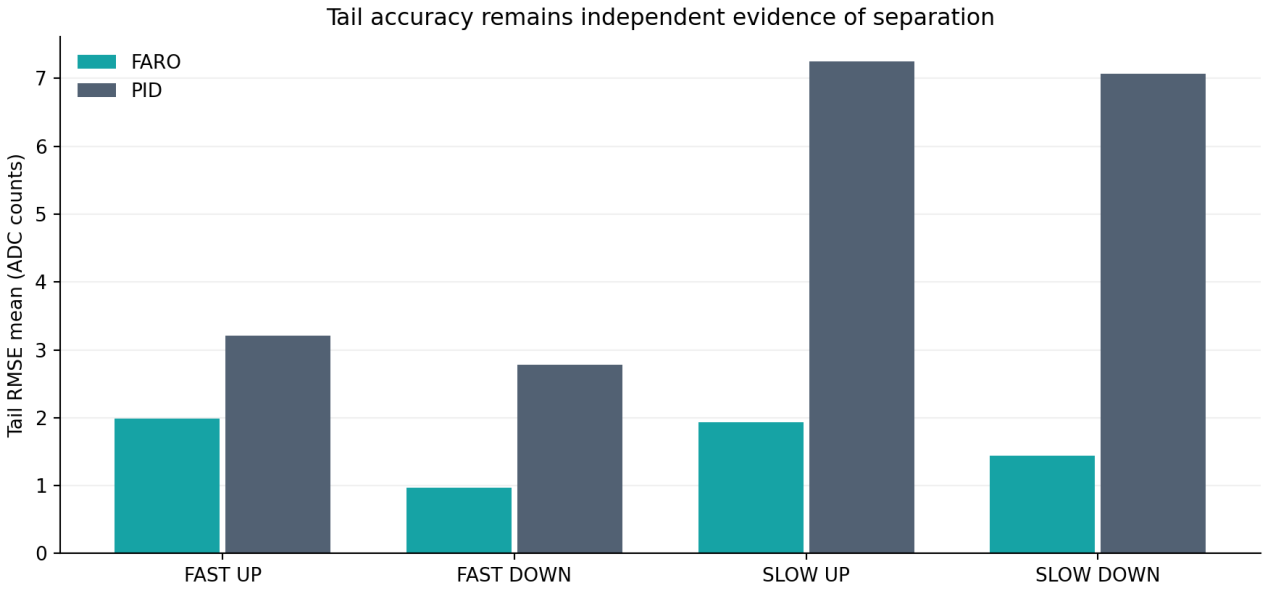
Profile	FARO IAE	PID IAE	Raw reduction	FARO duration	PID duration
FAST_UP	28,551.12	28,819.44	0.93%	65.502s	65.256s
FAST_DOWN	28,355.28	28,746.36	1.36%	65.519s	65.255s
SLOW_UP	29,094.36	29,408.88	1.07%	65.502s	65.261s
SLOW_DOWN	28,818.24	29,371.20	1.88%	65.503s	65.265s
Aggregate	114,819.00	116,345.88	1.31%	--	--

FARO recorded lower raw nominal IAE in FAST_UP, FAST_DOWN, SLOW_UP and SLOW_DOWN. All completed final-profile runs report zero cycles above 150 ms. The combined result file marks timing_comparable=true, while separately marking physical_comparative_claim_allowed=false because FAST and SLOW were separated by a contact intervention.

That distinction is preserved here. The result is evidence from a physical RC plant, but the four-profile aggregate is a descriptive consolidation rather than one uninterrupted mutation experiment.

6. Tail accuracy: a second view of separation

Tail RMSE is not affected by subtracting the IAE slew floor. It looks at the final 30 decisions of each stage and therefore emphasizes settled tracking rather than the compulsory initial ramp.

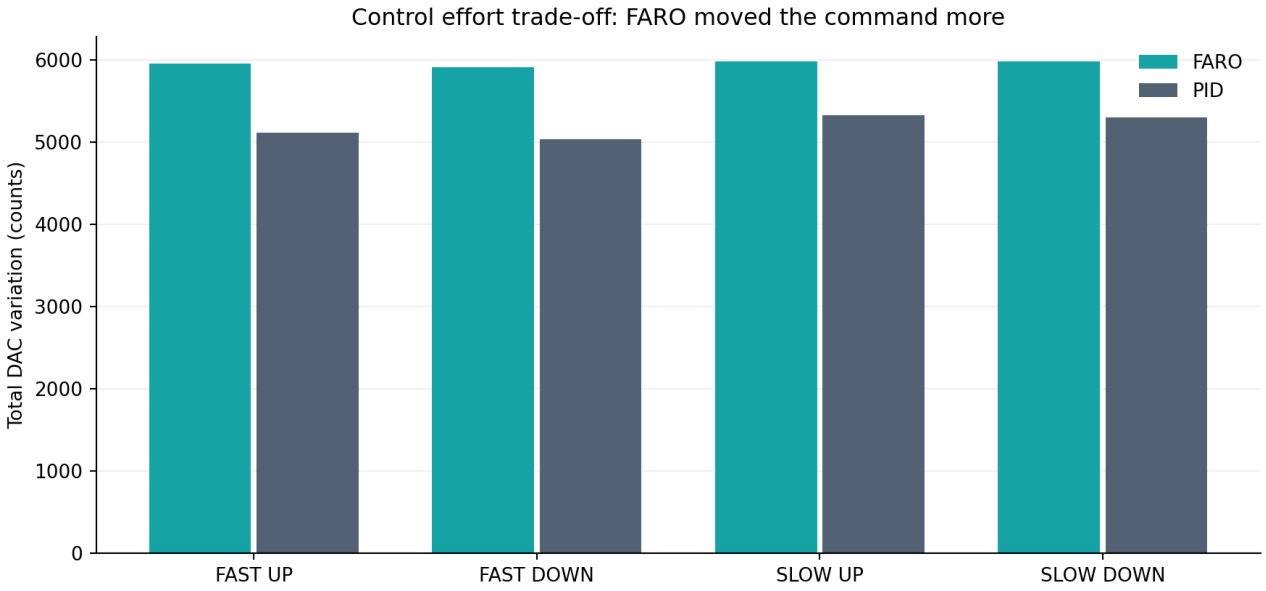


Profile	FARO tail RMSE	PID tail RMSE	FARO lower
FAST_UP	1.984	3.208	38.16%
FAST_DOWN	0.966	2.779	65.24%
SLOW_UP	1.937	7.253	73.30%
SLOW_DOWN	1.437	7.074	79.69%
Mean	1.581	5.079	68.87%

Across profiles, mean tail RMSE is 1.581 for FARO and 5.079 for PID, a 68.87% reduction. This supports the interpretation that the small raw aggregate IAE gap is partly masked by shared transient burden.

7. Control effort trade-off

FARO achieved lower tracking error with higher command movement. This trade-off must be shown alongside accuracy so the benchmark does not imply a free improvement.



Profile	FARO DAC variation	PID DAC variation	FARO vs PID
FAST_UP	5,955	5,115	16.42%
FAST_DOWN	5,909	5,033	17.41%
SLOW_UP	5,981	5,328	12.26%
SLOW_DOWN	5,981	5,294	12.98%
Total	23,826	20,770	14.71%

FARO total DAC variation is 14.71% higher than PID in this dataset. The report does not translate this count variation into power, wear, energy or actuator stress without a calibrated actuator model.

8. Methodological status and claim boundary

Status	Statement
Supported	Four completed profiles show lower raw nominal IAE for FARO than for the fixed R6 PID.
Supported	The exact slew-only causal floor implied by the 12-count action limit and reference schedule is 24,435.36 count-s per profile.
Supported	After subtracting that same floor from both arms, FARO is 8.21% lower aggregate on the above-floor remainder and 6.12%-11.20% lower by profile.
Supported	Tail RMSE is lower for FARO in all four profiles, while DAC variation is higher.
Not supported	Calling the entire above-floor remainder “100% controllable error.” RC dynamics can contribute unavoidable error beyond the command-slew floor.
Not supported	Claiming a fixed 86% floor. The exact fraction depends on which raw total is used as denominator.
Not supported	Claiming the combined FAST+SLOW aggregate is one uninterrupted physical mutation experiment. A contact intervention separates the blocks.
Not supported	Claiming universal superiority over PID, PID Oracle, RLS, MPC, LQR or other controllers from this benchmark alone.

The corrected interpretation is stronger than the raw 1.31% headline, but also narrower: it says the protocol itself compresses relative IAE separation, and that FARO performed better on the residual above a precisely defined common slew floor. It does not transform a post-hoc normalization into a preregistered primary endpoint.

9. Reproducibility record

The final SLOW continuation preserved the original firmware, FARO policy/update schedule, PID gains, profiles, 120 ms nominal period, scoring, qualification gate and initial-condition rules; it reused the MEDIUM prepared states and FAST results without repeating training or FAST.

Artifact	SHA256
Combined descriptive result	6196c6c7f2d912a81b3a4f4a4a85f0f862e03c27c98dfd1a7e2d5f44ceb05978
Continuation protocol	93460aedfcf3675c5fc6d6a88c9c8d5d9ac33afccc32b38d66dc6f24475b4d76
Continuation HASHES.json	7f6b5ae732ea6d84ee27713cc786e5d37bd7f8a2bb7f79ceb2f2f29095b4d5ff
Firmware HEX	75c1277b7a44e5c2ba379db1755f77c25b69f41b3704c0ce98747c142f46e9bd
OAPA helper DLL	07256634ec958d4ea343ae33cc40d8c74b35cf7ff029f310d0c121641d431a51
FARO OAPA V11 source	23b4ff1d9f24d763f396154c609b22aa90a9dfbff79c16722af459a69e62d4c2

Machine-readable companion files generated with this revision: SLEW_FLOOR_ANALYSIS.json and SLEW_FLOOR_RESULTS.csv. They contain the exact derivation, assumptions and recalculated percentages used in this PDF.

The benchmark retains the original contact-intervention disclosure. The next physical tests should be separate experiments: FARO vs PID Oracle; relay-driven in-run mutation; and FARO vs RLS. None of those future comparisons are implied by the present result.

10. Conclusion

Primary raw result: FARO 114,819.00 vs PID 116,345.88 nominal IAE - **1.31% lower**.

Exact slew-only floor: 24,435.36 per profile, 97,741.44 across four profiles - **84.01% of PID aggregate raw IAE**.

Above-floor residual: FARO 17,077.56 vs PID 18,604.44 - **8.21% lower** aggregate; per-profile reduction **6.12% to 11.20%**.

Settled tracking: mean tail RMSE 1.581 vs 5.079 - **68.87% lower with FARO**.

Trade-off: FARO command variation is **14.71% higher**.

Recommended public wording

“In the physical RC benchmark, FARO recorded 1.31% lower raw nominal IAE than the fixed R6 PID across four completed profiles. An exact post-hoc decomposition shows that the 12-count/decision actuator slew limit contributes a shared 24,435.36 count-s floor per profile. After subtracting only that common slew-only floor, FARO’s above-floor residual IAE is 8.21% lower aggregate (6.12%-11.20% by profile). Tail RMSE is 68.87% lower on average, while FARO uses 14.71% more total DAC variation. The FAST and SLOW blocks were separated by a contact intervention, so the aggregate is descriptive rather than one uninterrupted physical experiment.”

This wording is deliberately narrower than “FARO is 8.21% better.” It states exactly what was measured, exactly what was analytically removed, and exactly what remains unproven.